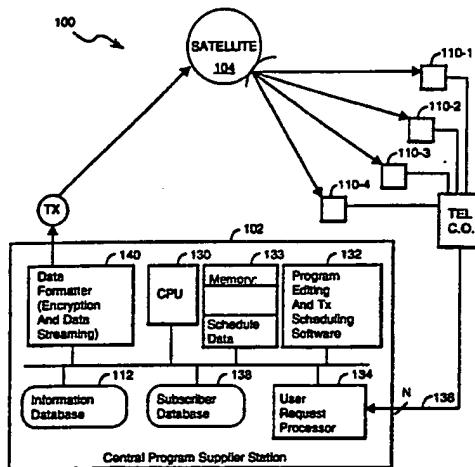




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : H04N 1/00, H04H 1/02, 1/00 H04B 7/185, G11C 7/00	A1	(11) International Publication Number: WO 93/09631 (43) International Publication Date: 13 May 1993 (13.05.93)
(21) International Application Number: PCT/US92/09087 (22) International Filing Date: 27 October 1992 (27.10.92)		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, SE).
(30) Priority data: 07/786,453 1 November 1991 (01.11.91) US		Published <i>With international search report.</i>
(71) Applicant: FINISAR CORPORATION [US/US]; 3515 Edison Way, Menlo Park, CA 94025 (US).		
(72) Inventor: LEVINSON, Frank, H. ; 3465 Cowper Street, Palo Alto, CA 94306 (US).		
(74) Agents: WILLIAMS, Gary, S. et al.; Flehr, Hohbach, Test, Albritton & Herbert, 4 Embarcadero Center, Suite 3400, San Francisco, CA 94111-4187 (US).		

(54) Title: INFORMATION BROADCASTING SYSTEM AND METHOD



(57) Abstract

An information broadcasting system (100) provides a large number of subscribers access to a large amount of information using one or more satellite transmission channels. The system (100) can also use cable television transmission channels. A program supplier station (102) stores an information database (112) and tags all the information in the database with indices so as to form a single hierarchical structure which encompasses the entire information database (112). Portions of the information database (112) are transmitted often, at least once per day, in order to provide the basic subscriber with information need to access the remainder of the database (112). The information provided by the basic subscriber service, which will typically include at least 50 gigabytes of data, is available to all subscribers without requiring two way communications between the subscribers and the program supplier station. Using a tiered system for scheduling transmission of the 50 gigabytes or so of information included in the basic subscriber service, as well as an intelligent subscriber request anticipation scheme for retrieving information before the subscriber asks for it, the present invention provides subscribers with reasonably quick access to all the contents of the large database while using only a modest amount of bandwidth.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NL	Netherlands
BE	Belgium	GN	Guinea	NO	Norway
BF	Burkina Faso	GR	Greece	NZ	New Zealand
BG	Bulgaria	HU	Hungary	PL	Poland
BJ	Benin	IE	Ireland	PT	Portugal
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SK	Slovak Republic
CI	Côte d'Ivoire	LJ	Liechtenstein	SN	Senegal
CM	Cameroon	LK	Sri Lanka	SU	Soviet Union
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	MC	Monaco	TG	Togo
DE	Germany	MG	Madagascar	UA	Ukraine
DK	Denmark	ML	Mali	US	United States of America
ES	Spain	MN	Mongolia	VN	Viet Nam
FI	Finland				

INFORMATION BROADCASTING SYSTEM AND METHOD

The present invention relates generally to the transmission or broadcasting of digital information to a wide base of subscribers, and particularly to methods and systems for providing access by many subscribers to an almost unlimited amount of information and/or programming material.

5

BACKGROUND OF THE INVENTION

10

The goal of computerized information servers is to provide a system that encourages use by subscribers, provides access to a large range of information, and which is flexible and inexpensive for information publishers. Computerized information services which exist today include Prodigy, Compuserve, and Dialog. Typically, the user either pays a monthly fee or a connect time fee for the most basic service and additional or special services are provided for additional fees.

15

20

All of the above mentioned computer services require two way communications between the user and the information server in order to serve each new request by the user. In other words, the user sends requests, and then selected data is transmitted individually to that user in response to his/her request. This two way "query based" communication scheme is also prevalent in other types of information servers. For instance, the "pay per view" services provided by cable television companies allow a subscriber to call the company and request a movie. The company then sends

- 2 -

5 a signal to the subscriber's television signal decoder to enable it to decode the program on a particular channel during the period of time associated with the requested movie. In this case, the "information" (a program or movie) is being transmitted in any case, but subscribers are only provided access to the program if they pay a special fee.

10 Two major distinctions between the present invention and the prior art query based information servers are (1) the present invention has a bandwidth that is thousands of times of greater than the prior art systems, thereby enabling high speed, low cost distribution of information, and (2) the present invention greatly reduces the amount of two-way communication required between subscribers and the information server by automatically transmitting information that responds to most anticipated "requests" by subscribers. Only those subscriber requests not satisfied by the automatically transmitted information need to be conveyed to the information server.

15 The goal of the present invention is to provide widespread, high speed access to a virtual omniscient database having typically well in excess of a terabyte (10^{12} bytes) of data. Using a data channel capable of transmitting, say, 1.5 megabytes of data per second, the present invention can provide virtually instantaneous access to about a 100 megabytes of information, high speed access (e.g., within two hours) to about a gigabyte of information, and can provide medium speed access, with perhaps 12 or 24 hour turn-around, to perhaps 100 terabytes of information. Furthermore, such access can be provided to a very large set of users without having to use a large number of data channels.

20 25 The large bandwidth of the present invention's information server enables the system to provide subscribers with multimedia programming, including video and audio programming. This is unlike current query based information servers, whose low bandwidth makes the distribution of video and audio programming impractical. The ability of the present invention to distribute

- 3 -

information in multimedia form makes it much more attractive to both subscribers and potential information publishers, including advertisers.

It is important to note that while a user has access to perhaps a terabyte, or even 100 terabytes or more, of data, the total amount of data that systems in accordance with the present invention system can transmit in any one day is much more limited, as will be described below. This is not unlike visiting the main library of a major university, such as Yale or Harvard University, having stacks containing several million volumes of books. Having "access" to all those books every single day does not mean that a user can receive them all in one day, nor does it mean that all the users can receive all the books in a single day. Nevertheless, each particular book (file or program) is available on relatively short notice, and having access to such a large collection of books (data) is still very useful.

The present invention also has tremendous cost advantages compared with information distribution using compact disks (CDs), and is believed to have distribution costs on the order of one thousand times less than any competing distribution media. The present invention eliminates both the need to generate "master" disks and the need to manufacture any physical media. Since distribution is accomplished by broadcasting, the actual costs for "publishing" a million copies of an average length novel (i.e., distributing it to a million subscribers) would be on the order of just a few dollars, versus a cost on the order a dollar or more per copy for printing either CDs or traditional books, plus similar costs for the physical distribution, resulting in traditional distribution costs totalling millions of dollars. Furthermore, updates to previously distributed information are also distributed at equally low cost using the present invention, in contrast with CD based publishing which requires mailing or otherwise physically producing and distributing new CDs to all previous purchasers.

SUMMARY OF THE INVENTION

In summary, the present invention is an information broadcasting system which provides a large number of subscribers access to a large amount of information using one or more satellite transmission channels. The 5 broadcasting system can also use cable television transmission channels or any similarly structured data distribution network. The system has a program supplier station which stores an information database and tags all the information in the database with indices so as to form a single hierarchical structure which encompasses the entire information database. Portions of the information database are transmitted often, at least once per day, in order 10 to provide the basic subscriber with information needed to access the remainder of the database. The information provided by the basic subscriber service, which will typically include at least 50 gigabytes of data, is available to all subscribers without requiring two way communications between the 15 subscribers and the program supplier station.

By using a "tiered" system for scheduling transmission of the 50 gigabytes or so of information included in the basic subscriber service, as well as an "intelligent" subscriber request anticipation scheme for retrieving information before the subscriber asks for it, the present invention provides a huge number 20 of subscribers with reasonably quick access to all the contents of the large database. This is accomplished even though only a modest amount of bandwidth is used. Furthermore, by reserving a portion of the system's bandwidth for satisfying requests for access to information not provided with the basic subscriber service, timely access to a virtually unlimited amount 25 of information can be provided, using the same modest transmission bandwidth, to thousands of those subscribers willing to pay additional fees for that service.

The utility of the present can be further enhanced by using data compression techniques so as to increase the bandwidth of the information server, and

- 5 -

by downloading software, data compression code books and the like along with the transmitted information. The present invention also includes the ability to automatically update the software used by subscribers' computers by transmitting the software updates along with other data transmitted by the information server, with different versions of the updated subscriber software being transmitted for each of the different platforms (i.e., types of CPUs) being used in subscribers' computers.

BRIEF DESCRIPTION OF THE DRAWINGS

Additional objects and features of the invention will be more readily apparent 10 from the following detailed description and appended claims when taken in conjunction with the drawings, in which:

Figure 1 is a block diagram of a "hyperbroadcast" data transmission system in accordance with the present invention, showing a receiving station suitable for use by a group of users.

15 Figure 2 is a block diagram of a typical receiving station for use by an individual.

Figure 3 depicts some of the data areas and software routines stored in the memory storage area of a subscriber station.

20 Figure 4 is a block diagram of a system for rebroadcasting onto a cable television network the data stream received from a satellite, as well as a local information source which can insert data into the data stream in lieu of data from the satellite.

Figure 5 is a block diagram of a local area network of subscribers.

Figure 6 is a block diagram of a data packet.

- 6 -

Figure 7 is a block diagram of a data filter subsystem used in subscriber stations for receiving broadcast data and video signals.

Figure 8 depicts the data structure of the filter list used to select data packets received by a subscriber station.

5

Figure 9 depicts the hierarchical structure of the broadcast data.

Figure 10 is a block diagram of a "user profile" data structure which specifies data to be stored and updated in a subscriber's computer.

Figure 11 depicts the index and timestamp information embedded in the broadcast data.

10

DESCRIPTION OF THE PREFERRED EMBODIMENT

15

Referring to Figure 1, there is shown an information broadcasting or transmission system 100 having a single program supplier station 102 which transmits a stream of data via a satellite 104 to a large set of receiving stations 110 owned or leased by subscribers 110. Subscribers are persons, organizations and companies with equipment to receive transmissions on one or more dedicated channels broadcast by the satellite 104.

20

In one implementation of the invention, it is expected that subscribers will have access to the basic service without charge. This basic service provides unlimited access to a base set of information, which in the preferred embodiment is about fifty gigabytes of public service information, including information such as newspapers, weather reports and the like that are updated frequently. Subscribers also have free access to all information on the service which advertisers or other suppliers of paid programming, such as governments and schools, have provided for distribution to the general user community. This access may include the ability to receive the broadcasted

25

- 7 -

5 advertising information as well as to make special telephone requests for additional information. Multiple channel versions of the system, which will include additional programming on separate transmission channels, will provide unlimited access to an even larger base set of information. In the preferred embodiment, all broadcast channels will include a mix of text, graphics, multimedia, audio, video and software programming, although the mix of programming materials will typically differ from channel to channel.

10 Subscribers who request transmission of data not included in the basic service could be charged for accessing that information. Information of this type is such that the provider of the information (e.g., a person or company owning copyrights on the information) has placed restrictions on how the information may be distributed. For example, the provider may require a transaction fee or royalty be paid before providing the subscriber with access codes to find or decrypt the information. Alternately, access to a set of information may be limited by a provider's requirement that he know who is receiving it. In either case, the subscriber, or his computer system, requests such information by calling the user request processor at a particular telephone number (toll free or otherwise) and specifying the information that the subscriber wishes to receive. The user is then given access codes and possibly decryption keys 15 to enable the user to access the requested information.

20

25 Alternately, subscribers could be charged a monthly fee for the basic subscriber service. Access by nonsubscribers, in systems which charge for the basic subscriber service, would be blocked either by encrypting the transmitted data, or by giving only subscribers the packet ID required to access the root information required to access the rest of the database, as will be described in more detail below.

The database 112 of information provided by the program supplier station 102 is indexed in a hierarchical fashion, providing keyed or indexed access to all data within the database 112. The database 112 will typically have

at least a hundred times more data that is available for retrieval than the data which is provided with the basic service. Subscriber requests for additional information are transmitted by conventional modems and telephone lines to the program supplier station 102, after which transmission of the requested information is scheduled for delivery (i.e., transmission) in the portion of the satellite's considerable bandwidth that is not used for transmitting the regularly scheduled basic programming. In the one-channel preferred embodiment, the system can transmit 32.4 gigabytes of specially requested data per day, which is expected to be sufficient to serve a subscriber base of approximately one million home users or several tens of thousands of commercial users. See Table 1 below. Additional transmission channels can be added as the subscriber base increases in size, and/or different satellite transmission channels can be used for different geographical areas.

As will next be described, the present invention uses a "tiered" system for scheduling transmission of the 50 gigabytes or so of information included in the basic subscriber service, as well as an "intelligent" subscriber request anticipation scheme for retrieving information before the subscriber asks for it. Together, these features of the present invention provide the system's subscribers with virtually instantaneous access to an large information database while using only a very modest amount of transmission bandwidth. Furthermore, by reserving a portion of the system's bandwidth for satisfying requests for access to information not provided with the basic subscriber service, timely access to a virtually unlimited amount of information can be provided to those subscribers, or advertisers, willing to pay additional fees for that service.

Overview of Basic Operation

Referring to Figures 1 and 2, the data stream is transmitted over a single satellite channel at a rate of 1.5 megabytes per second in fixed size 10KB (10 kilobyte) data packets having a maximum size of 10,000 bytes. The subscriber receiving stations 110 all include a data filter subsystem 120 (see

5

Figure 2) which receives the incoming data stream and selects those data packets that meet selection criteria defined by the receiving station 110. The information in the selected data packets is then downloaded into the memory 122-124 of the subscriber's workstation or other computer using software provided by the information network supplier for that purpose.

10

Typically, all subscriber stations are programmed to store a set of "root information", comprising about 250KB of information, which provides an index to the network's information database 112. Data packets containing this root information are marked to indicate the last time the root information was updated, and the subscriber stations are programmed to automatically replace their internal copy of the root information as updated information is received.

15

20

25

In addition, each subscriber can specify additional information from the database 112 that the subscriber wants to receive. All information in the system is hierarchically arranged, with the root information comprising the top several layers of the hierarchy. As will be described below, subscriber "requests" are made by selecting items from the top level menus, or by selecting "markers" in previously retrieved data that reference other portions of the network database. Each such marker includes a packet ID that is used by the receiver to select packets to be downloaded to the subscriber's computer, as well as a timestamp value indicating (A) whether or not the requested information is included in the basic subscriber service, and, if so, (B) the approximate time that the requested data is scheduled for transmission. Timestamps are described in more detail below. If the requested information is already scheduled for transmission sometime in the near future (either because it is part of the basic service, or because another subscriber recently requested the same item), the subscriber is informed of the expected arrival time.

If the requested data is not included in the basic subscriber service, and if the subscriber is authorized to request additional information (i.e., has signed

- 10 -

an agreement to pay for the transmission of addition data), the request for additional information is transmitted by conventional modems and telephone lines to the program supplier station 102, after which transmission of the requested information is scheduled for delivery (i.e., transmission) by the program supplier station. If the packet IDs associated with the requested information are not already included in the root information available to the subscriber, the program supplier station 102 also provides the subscriber with the packet IDs required to access the requested information. For information having an associated delivery charge, the subscriber's account is debited by a charge corresponding to the information requested.

In the preferred embodiment, each request is assigned a priority level (e.g., low priority night-time delivery, regular two-hour delivery, and rush half-hour delivery), and the program supplier station 102 schedules delivery of requested information accordingly. Information access charges are based both on the volume of information requested and on the request's assigned priority level. If the requested priority level is not available due to an unexpected surge in requests, the subscriber station is informed that the request priority has been downgraded and the subscriber is given the opportunity to cancel the request.

Program Supplier Station

As shown in Figure 1, the program supplier station 102 includes the aforementioned database 112, which is typically stored on a large number of high density magnetic disk devices. The program editing software 132 on the central program supplier station 102 tags all the information in the database 112 with indices (each of which contains a packet ID plus additional information) so as to form a single hierarchical structure that encompasses the entire information database. More specifically, software 132 (executed by CPU 130) generates a hierarchical set of indices referencing all the data in the information database 112 and embeds those indices in the information database.

- 11 -

5 In an alternate embodiment, the packet IDs and the associated indices embedded in the database can be generated off line. In particular, each information provider can be assigned, in advance, a block of packet ID values. If the packet ID values are made sufficiently long, such as 64 or 96 bits long, then assigning blocks of packet ID values to information providers will not deplete the supply of such values. Each information provider is required to organize the information it wants to publish into one or more subtrees, embedding the associated indices therein, and then providing the information to be published to an information server. The information server will then 10 "graft" the subtrees of information provided by each information provider onto the information hierarchy broadcast by the information server.

15 The program supplier station 102 includes at least one central processing unit (CPU) 130, software 132 executed by the CPU 130 for editing program materials and for scheduling delivery of both basic service information and subscriber requested data. In some embodiments, the supplier station 102 will be a distributed system utilizing many CPUs interconnected using a local area or even a wide area network. In that case, the station 102 shown in Figure 1 represents the information collection node, which collects all the information from various programming material providers, combines, serializes 20 and transmits the information.

25 Memory 133 stores, among other things, scheduling data which is used to control the timing at which each packet of information in the information database 112 is to be transmitted. As will be described below, the timing information in the schedule data is incorporated into "timestamp markers" that are transmitted along with the transmitted data, enabling subscriber stations to know in advance when selected information items will be received.

User request processor 134 is interfaced to a large number of conventional telephone lines 136, using individual modular interface cards for each telephone line so that the system is expandable. Each interface card is

- 12 -

5

10

programmed to receive and confirm information requests from subscribers, as well as to handle subscriber account transactions such as setting up new subscriber accounts, providing data decoding keys to new subscribers, and accepting credit card and debit card payments of subscriber balances. In response to each information request received, the request processor 134 queries a subscriber database 138 to ensure that the subscriber is authorized to request additional data, sends the request to the scheduling software 132 so that the transmission of the requested information will be scheduled, and enters the corresponding charge into the subscriber's records in the subscriber database 138.

15

In embodiments of the invention using multiple transmission channels, so as to be able to handle larger numbers of subscriber requests, special scheduling messages are transmitted via the satellite indicating the approximate time that the requested information will be transmitted and the transmission channel that will be used. Subscriber stations waiting for requested information use this information to change the transmission channel being monitored so as to receive the requested information.

20

The actual data transmission is handled by a data formatting and transmission subsystem 140. This subsystem includes memory for buffering and formatting the data scheduled for transmission, as well as circuitry for encrypting, packetizing, and then serializing or streaming the data to a transmitter 142 which transmits the data stream to the subscriber stations 110 via satellite 104.

25

In the preferred embodiment, some of the transmitted data is encoded using conventional public key encryption methodology. Generally, decoding keys for decoding the encrypted data are given only to subscribers who pay for access to that data. For instance, whenever a subscriber requests information not included in the basic subscriber service, and for which there is an associated delivery charge, the program supplier station 102 may select a

- 13 -

new encryption key for each request in a pseudo-random fashion, passing the corresponding decoding key to the requestor via the same telephone connection on which the request is received. Encryption and decoding keys are used in the preferred embodiment solely in conjunction with the distribution of materials published on a fee-per-purchaser basis. When multiple subscribers request the same set of materials prior to their transmission, the same encryption and decoding key can be used for all of those subscribers, thereby eliminating the need to retransmit the requested materials for each requesting subscriber.

Another type of "encryption" used in the preferred embodiment to defeat access by unauthorized subscribers to "protected data", without actually encrypting the protected data is as follows. The packet IDs associated with the protected data need not be sequential, rather they should be assigned on a pseudo-random basis. Furthermore, the packets for the requested data should be scheduled for transmission in a pseudo-random sequence, interleaved with other broadcast data packets. This method of "mixing up" the requested data, combined with the fact that at least a hundred 10KB packets are transmitted per second, will make unauthorized access to protected data extremely difficult even when the protected data is not encrypted prior to transmission.

The utility of the present invention can be further enhanced by using data compression techniques so as to increase the perceived bandwidth of the information server, and by downloading software, data compression codebooks and the like along with the transmitted information. Data compression is discussed more below.

Subscriber Receiving Stations

Figure 2 shows an individual subscriber station 110-1 having its own satellite dish 150, low noise block receiver and signal converter (LNB) 152, receiver 153, and data filter subsystem 120. Typically, the satellite dish receives a

- 14 -

5 signal having a carrier frequency on the order of 12 GHz, the LNB 152 brings this down to about 1 GHz with a signal bandwidth of about 500 MHz, and the receiver 153 then translates the received signal into a digital baseband signal. In this embodiment, the subscriber station 110-1 is a desktop computer (such as a Macintosh computer made by Apple Computer, or an IBM PC compatible computer made by any one of a large number of manufacturers) to which the satellite receiver 150-152-153 and data conversion subsystem 120 have been added. Thus, the subscriber station 110 has a CPU 154, monitor 156, keyboard 158, telephone modem 159 and other conventional computer components.

10

15 As noted above, the data filter subsystem 120 receives the incoming data stream and selects those data packets that meet selection criteria defined by the receiving station 110. The information in the selected data packets is downloaded into the memory 122-124 of the subscriber's workstation or other computer using subscriber software 160 provided by the information network supplier for that purpose. The data filter subsystem 120 is discussed in more detail below with reference to Figure 7.

20 The present invention also includes the ability to automatically update the software used by subscribers' computers by transmitting the software updates along with other data transmitted by the information server, with different versions of the updated subscriber software being transmitted for each of the different platforms (i.e., types of CPUs) being used in subscribers' computers. To implement this, the subscriber stations 110 must be set up to automatically accept and download software updates of the subscriber 25 software 160.

Referring to Figure 3, the subscriber's data storage 124 in the preferred embodiment includes a data area 161 reserved for downloading information from the data filter subsystem 120. Depending on the memory capacity of the subscriber's data storage devices 124, this data area 161 might be as

- 15 -

small as, perhaps, 2MB, or as large as perhaps 100MB, with a reserved data area of 10MB being more typical for an individual subscriber.

One feature of the present invention is that the subscriber software 160 has a set of subscriber request handling routines 162 which interact with the data filter subsystem 120 to arrange for downloading of the information requested by a subscriber. The request handling routines include a special "smart caching" routine 163 for anticipating future information requests by the user. Data corresponding to these anticipated information requests is downloaded into a portion 164 of the reserved data area 161 called the smart cache. Since a virtually unlimited amount of disk space could potentially be occupied by these anticipated requests, only some of which will actually be requested by the subscriber, the smart cache 164 is used as a temporary buffer, the oldest contents of which will be overwritten as information from new anticipated requests is downloaded.

The subscriber software 160 also includes a data decompression routine 165 for decoding information transmitted in compressed form, a cache management routine 166 for controlling the downloading of information into the data area 161 and determining which information in the data area 161 to delete when the area overflows, and a user profile setup program 167 which compiles a "user profile" data structure 300 that defines a default set of information to be retrieved from the transmitted data and stored in the data area 161. As noted previously, the data decompression routine used can change depending on the nature or type of data.

Figures 4 and 5 show two alternate configurations of subscriber stations, both of which reduce the cost per subscriber for the required receiver hardware. In Figure 4, the received satellite signal is translated by a low noise block converter (LNB) 152 to an intermediate carrier frequency of perhaps 1 GHz, which is then translated down to baseband by a receiver 153. The receiver 153 outputs a logic level bit stream that is received by a cable television

- 16 -

(CATV) distribution plant's head end 170 and then retransmitted via cable 172 to all the subscribers of the cable television system. Subscribers to the information network connect their computers to the cable television cable 172 using a conventional tuner 174 to bring the satellite signal down to baseband. The remaining portion of each subscriber's station is the same as shown in Figure 2. Two advantages of this system configuration are (1) that individual subscribers save the cost of having a satellite receiver 150-152, since one such receiver is sufficient to service a large number of subscribers, and (2) better reception of the satellite signal, with lower error rates, is possible since most CATV head ends have larger receiving dishes than can be afforded by individual subscribers.

An important feature of this configuration is data switch 176, which enables local programming to be added to the data stream received from the satellite. It is anticipated that a certain amount of the satellite's bandwidth will either be left unused, for instance by leaving room for 1000 locally generated data packets each time that a particular special packet is transmitted, or that a portion of the data stream will be designated as suitable for being preempted by local programming. The local program supplier station 178 is essentially a small version of central program supplier station 102 in that it must perform similar tasks - collecting information from various (presumably local) programming material providers, combining, serializing and transmitting the information during the available time slots set aside for local programming. The local programming data will also be hierarchically arranged with embedded timestamped indices, and will generally be set up as a subtree appended to a branch of the data hierarchy transmitted by the central program supplier station.

The ability to have local programming is commercially important, both because it provides the ability for local CATV stations to collect advertising revenues, which enables the data transmission service to be provided free to subscribers, and also to make the content of the transmitted data better tailored to its

audience. Vastly greater information content is another result of having multiple local source transmitters. By way of example, if the local program supplier for each local CATV system provides one percent of the total traffic, and if there are one thousand such local program suppliers, the overall information content of the information distribution system is increased ten fold.

In Figure 5, a single satellite receiver 150-152 is connected to a local area network 180 of subscribers. All requests for transmitted information are handled by a single data filter subsystem 120 and the network server 182.

In other words, requests for information are routed through the network server 182 so that one data filter subsystem 120 can be used to select all the information needed by all the subscriber work stations 184 on the local area network 180. Similarly, a single modem 186 or a small number of such modems can be used by a large number of work stations 184 to submit requests to the program supplier station. A significant cost saving associated with this configuration is saved disk space - because all of the users can share access to information stored on the server's hard disk. Speed of access to data from the information service may also be improved, compared to an individual subscriber configuration, because more disk space can be dedicated to automatically caching a larger portion of the transmitted information, since the cost of the dedicated disk space is being shared by a group of users.

This system configuration also spreads the cost of the satellite receiver 150-152 and the cost of the data filter subsystem 120 over the same group of users. Since it is anticipated that the data filter subsystem 120 might sell for at least \$500 (in 1991 U.S. dollars) apiece to manufacture, the local area network configuration of Figure 5 can result in significant savings per user.

The system configuration of Figure 5 is likely to be suitable not only for business subscribers, but also libraries, schools, kiosks at shopping malls, and other types of institutional subscribers.

Data Filter Subsystem

5

As shown in Figure 6, each received data packet 200 begins with a packet identifier 202, followed by a function code 204 and a set of data 206, followed by error detection codes 208. Furthermore, as will be discussed below, the data 206 includes not only the information associated with the packet identifier 202, but also timestamped indices 210 for the next N (e.g., three) levels of information in the information database hierarchy.

10

Packet identifiers 202 are used to determine which packets require further processing, with all unselected packets simply being discarded by the subsystem 120. The function code 204 is used to identify and categorize special messages transmitted by the program supplier. The packet ID 202 and the function 204 are not encrypted, thereby allowing packets to be selected or discarded prior to processing the data 206. In the preferred embodiment, the data 206 portion of most packets is not encrypted, except for the use of data compression. A standard, commercially available error correction methodology may be used to detect and correct data errors caused by noise in the transmission process.

15

In some applications, data packets transmitted in response to subscriber requests will be encrypted. However, given the huge amount of data being transmitted it is impractical for a "pirate" to scan all the transmitted data for useful information. Without knowing the packet IDs for a particular set of information, and the proper order of those packet IDs for reconstructing a particular file or program, it is virtually impossible to extract useful information from the data stream, thereby making data encryption unnecessary in most cases.

20

Most or all of the data transmitted by the system will be transmitted in compressed form. In order to make the data compression both efficient and flexible, a general set of data decompression software is included in the subscriber software 160 (see Figure 2) given to each subscriber, along with

- 19 -

5 a default "code book" or coding algorithm to be used for the decompression process. Other code books or coding algorithms, associated with various types of transmitted data will be transmitted along with the compressed data. A typical code book, for example suitable for encoding newspaper articles, will occupy less than 50KB. When used in conjunction with megabytes of data, and taking into consideration that the compression will typically allow two-to-one or better data compression, the transmission of code-books along with the related data involves an inconsequential amount of overhead. The data decompression is performed on the subscriber's own computer, where 10 computation time is presumably virtually free of cost. Also, data decompression is not time consuming computationally using currently available desktop computers, and will not noticeably impact the performance of the system from the viewpoint of subscribers.

15 Referring to Figure 7 the data filter subsystem 120 used in subscriber stations to receive broadcast data works as follows. The satellite transmitted data stream, which contains a sequence of data packets, is first processed by a conventional satellite signal recovery circuit 220 that stores the received data packets in a ring buffer 222. In the preferred embodiment, the ring buffer 222 has 4 megabytes of storage, enabling about four hundred 10KB data 20 packets or about 2.5 seconds data at the system transmission rate to be stored. The required size of the ; buffer 222 is governed by the maximum possible delay before a selected packet is downloaded onto the subscriber's host computer.

25 Each time that a complete data packet is transferred by the data recovery circuit 220 to ring buffer 222, the data recovery circuit 220 sends a "packet received" message to the subsystem's controller 224. The controller 224 responds by comparing the packet's ID 202 with a list of packet ID values and/or ranges of packet ID values stored in filter list 226. Packet's which match any of the packet ID values in the filter list 226 are "selected" and

- 20 -

downloaded to the subscriber's computer via a conventional bus interface circuit 234, such as a SCSI (Standard Computer System Interface) interface.

5

To stretch the time available for downloading, a modified ring buffer technique can be used in which the buffer 222 is divided into a predefined number of slots, each of which can hold one packet. Slots holding packets selected for downloading are "marked" (e.g., using a set of bit flags) and thereby prevented from being overwritten by incoming data until the packet in the slot has been downloaded and the slot unmarked.

10

If a substantially longer period of time is required to ensure successful downloading of the received data to the subscriber's computer, either additional random access memory can be added to the ring buffer 222, or a fast disk storage device can be added to the data filter subsystem 120 to enable larger amounts of data to be buffered by the subsystem 120. The use of disk storage device for buffering may be justified for local area networks of subscribers, primarily to avoid loss of data when the network server cannot service the subsystem 120 within a few seconds.

15

20

In embodiments of the invention which include the transmission of video or audio programming material, the video programming material is transmitted digitally, in a compressed form. Many video data compression systems are commercially available. The video program material stored in the information database 112 is stored in compressed digital form by the program supplier station 102. As a result, each data conversion subsystem could optionally include a video decompression and modulation circuit 240 for decompressing the compressed digitized video data, and then converting it back into an analog signal suitable for either playback or recording using a conventional VHS video recorder. In alternate embodiments, the received video data will be downloaded in compressed form, for example for storage on a digital tape recorder. In this alternate embodiment, the subscriber's computer will have to include a video decompression circuit. This embodiment has the advantage

25

- 21 -

are reducing the storage requirement associated with digital video signal storage.

When multiple transmission channels are used by the program supplier station, the controller 224 sends control signals to the satellite receiver as to which channel should be selected. Typically, the satellite receiver will normally be tuned to a "main" or primary channel. Other channels are only selected during the expected transmission times associated with packet identifiers stored in the filter list.

Referring to Figure 8, the "requested data" in the filter list 226 is stored as an ordinary B-tree data structure, similar to those found in almost any DBMS (database management system). At the top of the tree is a root node 260. At the bottom of the tree are leaf nodes 266, 268. In between the root 260 and the leaf nodes are intermediate nodes 270, 272, 280. B-trees are well known data structures used to quickly identify whether or not a specified key value is present in a database. In this case, a three level tree, with perhaps up to twenty entries per node, will facilitate high speed determination as to whether or not a particular packet ID is present in the filter list.

Each leaf node in the filter list contains five fields: (1) a packet ID, (2) a flag indicating whether the packet is provided as part of the basic subscriber service or whether it is a packet that has been specifically requested to be transmitted, (3) the identify of the channel on which the packet is to be transmitted, (4) a repetition rate value, which is the rate at which transmission of the packet is repeated (e.g., one time per hour, or zero for information specifically requested), and (5) a time skew value indicating the approximate time that the packet is scheduled to be transmitted within the packet's repetition tier.

To determine the approximate next time that a packet is will be transmitted, the repetition rate and time skew fields are used:

- 22 -

... Period + INT{Current Time/Ran Period}